Investigating the Impact of Cyber Security Attacks on Cryptocurrency Markets

A thesis submitted in partial fulfilment of the requirement
for the degree of Master of Research

Submitted by:

Seung Ah Lee

Supervised by:

Associate Professor George Milunovich *and* Dr Colin Zhang

August 10, 2022

Department of Actuarial Studies and Business Analytics
Macquarie Business School
Macquarie University

Statement of Originality

This work has not previously been submitted for a degree or diploma in any university. To the best of my knowledge and belief, the thesis contains no material previously published or written by another person except where due reference is made in the thesis itself.

Seung Ah Lee
June 16, 2022

Acknowledgements

This thesis has been completed with the support from many people that I would like to acknowledge.

Firstly, special thanks to my supervisor George Milunovich. I would not have been able to complete this thesis without his guidance and help. Whenever I faced a difficult problem, his encouragement, insights and knowledge inspired me and became a driving force to keep moving forward. He is the best teacher and a role model I want to emulate the most.

I would also like to thank Colin Zhang, my associate supervisor. Several years after my coursework studies he graciously accepted my request to act as my co-supervisor. He has been very supportive, and his feedback is always constructive.

Lastly, my dear family and friends. No words of gratitude are enough for their love and support. My parents and big brother give me endless love and unconditional support. My big sister and mentor, Charlotte Park, has always been by my side and has been with me through tough times. Your encouragement, support and care give me so much strength and confidence. Bo Sun Kim, who has always been my best friend, encourages me, gives me strength in difficult times, and believes that I will get through it no matter what. Daseul Baek, a friend who gives me strength more than anyone else and encourages me to keep moving forward. Thank you, and love you all. I will never forget the support and strength you gave me.

Abstract

Cryptocurrency markets have grown significantly since the introduction of Bitcoin in 2008. Currently, there are more than 500 cryptocurrency exchanges worldwide and over 19,700 different cryptocurrencies which trade across various markets. While cryptocurrency trading is possible via peer-to-peer transactions, more than 90 percent of trading occurs on organised exchanges. Therefore, centralised cryptocurrency exchanges have become high-value targets to hackers and other types of criminal activity. In this thesis, I investigate two aspects of risk associated with cyberattacks on digital exchanges. First, I study the risk of cryptocurrency exchange closures and attempt to predict which markets will remain active given publicly available data on their key characteristics, including cybersecurity measures. I construct predictive models which reach training set accuracy of up to 95.9 percent, and up to 85.7 percent accuracy when applied to independent test data. In terms of feature importance, I find that transaction volume, exchange lifetime and cyber security measures such as security audit, cold storage and bug bounty programs rank high in their contribution to the predictability of exchange closures. Second, I examine the impact of cybersecurity breaches of cryptocurrency exchanges on the return of Bitcoin. Using several alternative specifications, I test the hypothesis that Bitcoin returns experience a decrease on the dates associated with cybersecurity breaches of digital markets. I find a negative and statistically significant impact at the 5 percent level, where Bitcoin price declines between 1.29 and 1.47 percent on cyberattack days, depending on which model is applied and what control variables are included.

Keywords: Cryptocurrency Exchanges, Closures, Digital Coin, Machine Learning, Predictive Analytics, GARCH, Return

Contents

List of Figures

List of Tables

1 Introduction

The phenomenon of cryptocurrency emerged in 2008 when Bitcoin was introduced as the first cryptographic currency that transacts on a peer-to-peer network (Nakamoto, 2008). According to a definition from the Reserve Bank of Australia (RBA, 2022), cryptocurrencies are digital tokens that allow people to make payments directly to each other through an online system without an intermediary such as a bank. Regarding their value the RBA states that "cryptocurrencies have no legislated or intrinsic value" and that they are worth what people are willing to pay for them in the market. Nevertheless, the cryptocurrency market has grown significantly since 2008. Currently, there are more than 500 cryptocurrency exchanges worldwide, and over 19,700 different cryptocurrencies trade across various markets. The market is, however, concentrated in the top 20 cryptocurrencies which account for over 90% of the overall market capitalisation of $1.88 trillion[1].

While the word cryptocurrency implies digital (virtual) currency secured by cryptography, there is considerable debate regarding whether cryptocurrencies are in fact currencies or if they should be classified as an asset class. For instance, Cheah and Fry (2015) highlight Bitcoin's speculative properties and suggest that Bitcoin is better characterised as an asset than an alternative currency. This view is supported by Glaser et al. (2014), who argue that Bitcoin should be considered a speculative financial asset due to its high volatility. Similar conclusions are provided in Bouri et al. (2017), Stensås et al. (2019) and others.

According to trading records the initial transaction volume of Bitcoin was minimal, and its price was close to zero USD. For instance, in one of the first Bitcoin transactions 10,000 Bitcoins were exchanged for two pizzas (Wallace, 2011). However, the concept of peer-to-peer payment system attracted sufficient interest, the number of transactions increased, and new cryptocurrencies were soon after created. Today, the value of 10,000 Bitcoins exceeds USD 416 million. Following the first few years of trading, the price of Bitcoin assumed a steep upward trajectory and reached its peak of USD 68,000 in November 2021. However, such price increases have also been accompanied by large volatility. Currently, the cryptocurrency market is still dominated by Bitcoin, which accounts for about 41% of the total market value[2].

There are several factors behind the extraordinary growth of cryptocurrencies. Their key characteristics may be listed as follows: i) decentralisation, ii) anonymity, iii) transactions which are irreversible and immutable, iv) security and v) fast and easy access. Each of these properties played an important role in the fast adoption of the new technology and warrants a brief explanation which is provided below in the context of Bitcoin, while the technology behind other cryptocurrencies may differ somewhat.

Unlike traditional fiat currencies, cryptocurrency is a decentralised virtual currency (Poletti, 2018). Here *decentralisation* refers to both the transfer of control, and to the decision-making regarding the supply of currency. Traditionally, fiat currency transactions are cleared by a third party within a centralised payment system run by a bank or some other institution. In contrast, decentralised payments are conducted directly between two counterparties (Luther and Smith,

[1] As of Jun, 2022. Source: https://coinmarketcap.com/.
[2] As of Jun, 2022. Source: https://www.tradingview.com/symbols/CRYPTOCAP-BTC.D/.

2020). A blockchain based system processes and validates transactions through an open network, and all transactions are recorded in a publicly available ledger called blockchain. Furthermore, the issuance of new coins, e.g. Bitcoins, is not under the control of government authority such as a central bank, but is instead managed by the rules embedded in the open source code that runs the network.

Cryptocurrencies can be stored in online 'hot wallets' or offline 'cold wallets' which are accessed using a private key (Guri, 2018) and provide *anonymity*. In the case of traditional fiat currencies, transactions require the identity of the remitter and recipient, while cryptocurrency transactions do not require a proof of identity. Anyone with access to a digital wallet and a private key can send and receive Bitcoins. The private key is a randomly generated string used to prove ownership and allows cryptocurrency to be spent. It is always mathematically related to a digital wallet address but is impossible to reverse engineer thanks to encryption which provides *security*. Next, cryptocurrency transfers conducted on the Bitcoin network are relatively *fast, cheap and efficient*, especially when compared to international transfers of fiat currencies (Masilela et al., 2021). Once a Bitcoin transaction has been processed, it is recorded on the blockchain and cannot be cancelled or modified, thus it becomes *irreversible and immutable* (Feig, 2018).

Despite the fact that cryptocurrency trading is possible via anonymous peer-to-peer transactions, most of it occurs on organised exchanges. In general, cryptocurrency markets can be characterised as centralised and decentralised. Centralised exchanges are intermediaries, most of which are incorporated as private companies that facilitate trading in cryptocurrencies. On the other hand, decentralised exchanges offer trading platforms where buyers and sellers decide the exchange rate and payment method (Matkovskyy, 2019). Nevertheless, it is estimated that about 99 percent of transactions are made through centralised markets (Roubini, 2018), and hence in this thesis I will limit the analysis to this type of cryptocurrency exchange. There are several reasons for the popularity of organised exchanges. First, Barbon and Ranaldo (2021) report that centralised exchanges have lower transaction costs and higher liquidity than decentralised exchanges. Another critical factor behind the popularity of centralised exchanges is that they simplify cryptocurrency trading by providing their clients with custodian services whereby traders can store their cryptocurrencies in accounts provided by the exchanges. This eliminates the need for setting up own digital wallets, which in some cases can be a technically challenging process.

While providing convenience, the practice of storing digital currencies on cryptocurrency exchanges presents a source of risk to cryptocurrency traders. Oosthoek and Doerr (2020) suggest that the lack of regulation in the cryptocurrency market and the fact that cryptocurrency prices have grown exponentially have made this market a high-value target for many types of cybercrime. Empirically, digital exchanges have a high incidence of being hacked or becoming a victim to some other form of a cybersecurity breach in which client funds are stolen from the exchange. From a cryptocurrency trader's point of view, this creates two types of risk. First, there is the risk of having investor funds stolen in a cryptocurrency exchange security breach. Second, cryptocurrency prices may be negatively impacted by successful cyberattacks on digital exchanges resulting in capital losses. For instance, the cybersecurity heist of Mt. Gox exchange

in 2014 led to a 36% drop in Bitcoin price (Hu et al., 2020). This can happen for several reasons. First, there is the possibility that a large amount of stolen cryptocurrency will be placed on the market resulting in a temporary oversupply. Another cause of such price reaction can be linked to the fact that many cryptocurrency exchanges faced closure following a cybersecurity breach thus undermining the infrastructure of the system. Lastly, as discussed in Kamiya et al. (2021), cyberattacks which involve the loss of personal financial information have large reputation costs that far exceed the out-of-pocket costs associated with the cyberattack itself.

As cryptocurrencies become a more mainstream form of investment and finance many market participants remain unaware of specific risks encountered in the cryptocurrency markets. In this thesis, I investigate the two types of risks associated with cryptocurrency exchange cybersecurity breaches mentioned above using cross-sectional data and time series data, namely i) the risk of potentially losing funds to digital market closures, and ii) the risk of cryptocurrency capital loss resulting from a price drop associated with successful cyberattacks on digital exchanges. In particular, I address the risk investors face due to closures of digital exchanges by attempting to forecast such closures on the basis of publicly available data. If it is possible to accurately predict which markets will remain open, and which ones will go out of business, then investors can account for this information and avoid exchanges that are likely to face closure. Here I show that it is possible to predict which cryptocurrency exchanges will face closure with relatively high accuracy using publicly available data. Next, I proceed to study the indirect effect of Bitcoin price reaction to cyberattacks on digital exchanges. I show that this is indeed a risk investors need to be aware of as successful cyberattacks on cryptocurrency markets have a significant and negative effect on Bitcoin price. Bitcoin price is found to be negative and range between -1.288% and -1.470%, depending on model specification, on the days of successful cyberattacks on cryptocurrency exchanges. The findings provided in this thesis can be taken into account by cryptocurrency market participants when formulating risk management policies.

This thesis consists of 6 chapters. Chapter 1 is the introduction and Chapter 2 provides a literature review. Data collection methods and the summary of the dataset are presented in Chapter 3, Chapter 4 outlines empirical techniques used in the thesis, while empirical results are contained in Chapter 5. Finally, Chapter 6 concludes and discusses some limitations of the study.

2 Literature Review

This literature review contains three subsections which cover three related topics on cryptocurrency markets. In the first subsection I discuss the literature which examines the impact of cybersecurity attacks on cryptocurrency exchanges. This is followed by an overview of the studies that investigate empirical characteristics of cryptocurrency returns and volatilities. Finally, I review the literature examining the impact of cyberattacks on cryptocurrency returns.

2.1 The Impact of Cybersecurity Attacks on Cryptocurrency Exchange Closures

Cryptocurrency exchanges typically operate as private companies in a largely unregulated market. Since a large majority of cryptocurrency transactions are conducted on centralised exchanges, which also often store currencies for their clients, exchanges are being targeted by cybercriminals as high-value targets. Consequently there has been a large number of digital exchange closures resulting in the loss of investor funds. In this section I review several studies that investigate links between cryptocurrency market closures and various exchange characteristics, and seek to uncover the factors behind the high incidence of cryptocurrency exchange closures.

Moore and Christin (2013) is one of the first papers to study the risk of cryptocurrency exchange closures. They examine 40 Bitcoin exchanges from 2011 to 2013 and employ survival analysis to identify which factors are behind Bitcoin exchange closures. Their model uses three variables to explain Bitcoin exchange closures, namely: i) average daily trading volume, ii) experiencing a previous security breach, and iii) Anti-Money-Laundering and Combating the Financing of Terrorism (AML/CFT) compliance. Using these variables the authors estimate a proportional hazards model and report that daily volume negatively impacts closures, i.e. the higher the trading volume the lower the exchange closure probability. They also use a logistic regression to examine which factors contribute to exchange cybersecurity breaches (without necessarily resulting in exchange closures) where the average daily transaction volume and months operational are the explanatory variables. The results show that transaction volume is the key factor behind cryptocurrency exchange cybersecurity breaches impacting the probability of experiencing a new breach with a positive coefficient, while the coefficient on months operational is negative. In summary, Moore and Christin (2013) find that cryptocurrency exchanges with high transaction volume are less likely to close down, but that they have a greater chance of being breached.

Moore et al. (2018) extend Moore and Christin (2013) to a longitudinal study over the 2010 — 2015 period. Their dataset includes the following variables: incidence of previous cyber breach, average daily trading volume, security properties of exchanges such as two-factor-authentication, bug-bounty program, security audit, and the implementation of cold storage. An additional predictor is the anti-money laundering and combating the financing of terrorism (AML/CFT) index for the home country of each exchange. In this study's data sample there are 25 exchanges that have been victims of hacking attacks (or other types of criminal activity), 15 of which subsequently closed and 10 markets that survived. In addition, another 23 exchanges closed without a breach experience. A critical aspect of exchange closures is whether they reimburse clients after the closure. In only 16 cases of all 38 exchanges that closed, customers were reimbursed either fully or partially. Moore et al. (2018) apply a methodology consisting of a logistic regression with fixed effects which is estimated using maximum likelihood estimation in R (pglm). In their baseline model, the authors include an intercept, breached in the current quarter variable, log of transaction volume and time trend. They report all estimated coefficients to be statistically significant at the 10 percent level, while breached in the quarter variable

is statistically significant at the 1 percent level. In addition, the coefficients on the log of transaction volume and the time trends are negative. At the same time, the breached in the quarter parameter is positive, indicating that experiencing a previous security breach in the quarter increases the probability of closure.

Another study that investigates cybersecurity breaches of digital exchanges but in the context of attack pattern is Oosthoek and Doerr (2020). Their methodology is based on attack vector analysis using the Vocabulary for Event Recording and Incident Sharing (VERIS) technology for post-breach assessment and impact analysis. VERIS is a Cyber Threat Intelligence (CTI) provided by Verizon and provides breach incidents in a structured form. CTI is used to analyse the tools, tactics, and procedures of cybersecurity breaches. The attack vector analysis yields four key findings. First, the increase in breach incidents reduces disclosure details. Authors suggest that when a breach incident is dated in the early years of the sample period, the details of the breach are typically unavailable due to the closure of those exchanges. However, the information is scarce even in more recent incidents because the exchanges are reluctant to share incident details in order to protect their reputation. Second, they find a decrease in the use of stolen credentials over time (in the early incidents most breaches were performed using stolen credentials). Third, there is a decrease in the abuse of functionality where attackers use legitimate methods to effectively abuse native functionality on an exchange. This type of breach is usually performed due to inadequate monitoring software or security audits of the exchanges. Finally, the authors find that in the case of Bitcoin exchanges, funds are stolen more often than in cyberbreaches of any other type of financial institution. However, in the case of recent breaches, investors have received either partial or full reimbursement from the affected exchanges and the amount of stolen BTC has decreased over time.

The final paper I review is Milunovich and Lee (2022) which is based on a preliminary version of this thesis. In that paper I use four classifiers: decision tree, random forest, logistic regression and support vector machine to predict which digital exchanges will close down and which ones will remain in business based on their publicly available attributes. The paper employs a database on 238 digital exchange which is collected manually from various media portals and internet searches. In contrast in this final version of the thesis I extend the number of predictive models from four to ten, and update the dataset to include 279 exchanges. Milunovich and Lee (2022) find that the four classifies they employ perform well, but that random forest outperforms the other models reaching training accuracy of 0.904 and 0.861 accuracy on independent test data. Critical features for predicting which exchanges will stay in business are: i) average trading volume, ii) lifetime of exchange, iii) security audit, iv) bug-bounty, and v) cold storage. In comparison, two-factor authentication and AML/CFT index do not appear to contribute much to the forecast accuracy.

2.2 Characteristics of Cryptocurrency Returns and Volatilities

The literature investigating empirical characteristics of the cryptocurrency markets is considerably larger than the above reviewed literature on cryptocurrency exchange closures, which is relatively new. Below I review a number of relevant papers that study first and second moment

dynamics of various cryptocurrencies.

Koutmos (2018) studies the relationship between Bitcoin returns and trading activity using a bivariate vector autoregression (VAR). The study is based on a dataset consisting of Bitcoin spot prices and transaction volumes over the period 2013 – 2017, comprising 1,231 observations. The author employs two different bivariate VAR models. The first model is specified for Bitcoin returns (percent changes in log price) and percent changes in the total number of Bitcoin transactions. The second model is Bitcoin percent returns and percent changes in the number of Bitcoin addresses. The author finds a significant relationship between Bitcoin returns and its transaction activity. Bitcoin returns typically increase by 0.3% three days following a simulated shock to trading activity.

The relationship between cryptocurrencies and macro-financial market risk factors is studied in Koutmos (2020). The study is conducted within the framework of Markov regime-switching regression analysis. In particular, the author investigates whether Bitcoin is a unique asset class that is not linked to economic fundamentals as represented by commonly used market factors. The study concludes asset pricing risk factors such as interest rates, implied stock market volatility and foreign exchange market variables are important determinants of Bitcoin returns. However, Bitcoin returns are more difficult to explain during periods of high volatility than during periods of low volatility. The effect of macroeconomic news announcements on Bitcoin returns is also reported in Pyo and Lee (2020).

Another study that investigates the relationship between Bitcoin and macro-financial factors is Van Wijk (2013) who finds significant long-run effect from the Dow Jones index, WTI oil and USD exchange rate to the price of Bitcoin. The finding that cryptocurrency markets are not isolated from macro-financial factors is also found in Bouri et al. (2018). They estimate STGARBTGARCH-M model and find spillover effects between Bitcoin, commodities, stocks, currencies and bonds, where the estimated links are stronger for returns than volatilities. Furthermore Bitcoin is typically found to be the recipient of the estimated spillovers, rather than a transmitter. Additional evidence against the hypothesis that cryptocurrency markets are isolated is found in Corbet et al. (2020b) where the Bitcoin market is found to respond to macroeconomic news and in Zhou (2021) who concludes that Bitcoin is likely to move with global financial markets at times of uncertainty. Similarly, Klein et al. (2018) applies BEKK-GARCH specification to six time series including Bitcoin, gold and silver prices in USD, crude oil prices for the West Texas Intermediate (WTI), the S&P 500 index, MSCI World and the MSGI Emerging Markets 50 index between July 2011 and December 2017. They report that Bitcoin returns show asymmetrical movements in response to market shocks in the same direction as precious metals. Moreover, they find Bitcoin returns to be positively correlated with the US equities during downward trending markets.

In contrast to the above mentioned studies a number of papers report that the cryptocurrency markets are largely isolated from mainstream asset class and do not depend on macro-financial risk factors. For instance Baek and Elbeck (2015) finds no relationship between Bitcoin and macro-financial factors using linear regression analysis. Ciaian et al. (2016) apply VAR analysis to find a significant impact of global macro-financial factors, as captured by the Dow Jones Index, exchange rate and oil price, on Bitcoin price only in the short run, but conclude that traditional

risk factors do not determine Bitcoin price in the long run. Finally, Briere et al. (2015) use weekly data over the 2010–2013 period to analyse a Bitcoin investment from the standpoint of a US investor with a diversified portfolio including both traditional assets (worldwide stocks, bonds, hard currencies) and alternative investments (commodities, hedge funds, real estate). Over the period under consideration, Bitcoin investment had highly distinctive features, including exceptionally high average return and volatility. Its correlation with other assets was remarkably low and provided significant diversification benefits.

Another strand of the literature shows that the volatility of Bitcoin returns exhibits time-varying dynamics. For example, Katsiampa (2017) studies the volatility characteristics of Bitcoin using several GARCH models. The dataset used spans the time period July 2010 – October 2016 and consists of 2267 observations. Bitcoin returns are modelled using an autoregressive (AR) model, and a first-order GARCH-type specification is employed in the conditional variance equation. According to the information criteria the AR(1)-CGARCH(1,1) (component GARCH) model appears to be optimal amongst several alternative specifications. The residual tests for this best model indicate no autocorrelation remaining in the residuals and that all of the conditional heteroskedasticity has been eliminated from the standardised residuals according to the ARCH(5) test. In another application of GARCH models Dyhrberg (2016) employs an asymmetric GARCH specification to demonstrate that Bitcoin may be useful in risk management and ideal for risk averse investors in anticipation of negative shocks to the market.

Bitcoin plays a key role in the cryptocurrency market among many cryptocurrencies with 46% market dominance. Several studies examine the relationship between Bitcoin and other major cryptocurrencies.

Kumar and Anandarao (2019) examine the return and volatility spillovers between four major cryptocurrencies using a GARCH model. They collect daily log returns of Bitcoin, ethereum, ripple and litecoin between August 2015 and January 2018 and estimate a DCC-IGARCH model. They find the existence of substantial spillover effects among all cryptocurrency returns. Moreover, the volatility spillovers between Bitcoin on one hand, and Ethereum and Litecoin on the other hand, are largest. The magnitude of the estimated volatility spillovers increases after 2017 as the trading activity in cryptocurrencies gains momentum.

Candila (2021) adopts a mixed-frequency approach within the Dynamic Conditional Correlation (DCC) specification in order to model the co-movement of seven key cryptocurrencies by adding Google queries data as a potential driver of volatility. The investigated time period is June 2016 to December 2020. This paper solves the issue of including low frequency (monthly) data on Google searches of each digital currency in the DCC framework modelling high frequency (daily) correlations between different cryptocurrencies. Candila (2021) reports that the inclusion of the monthly Google searches as additional determinants for the daily volatilities of the chosen digital currencies is important. They find that only the models using Google searches belong to the set of superior models (SSM), identified through the Model Confidence Set (Hansen et al., 2011) procedure. Moreover, only the models employing the Google trends and the RM model have satisfactory residual diagnostics. Finally, the estimated time-varying correlations are relatively high, ranging between 0.6 and 0.8 on more recent data.

Finally, Gradojevic and Tsiakas (2021) study volatility transmissions between long and short

time horizons and whether such transmissions depend on the type of volatility, i.e. low vs. high volatility, in the context of three cryptocurrencies: Bitcoin, etherium and ripple. For instance, one of the investigated questions involves examining whether high long-horizon volatility will result in high short-horizon volatility and vice versa. Using a method based on a wavelet Hidden Markov Tree model specifies a two-state regime of high and low volatility. Their main empirical finding is that going from long to short horizons, volatility cascades tend to be mostly symmetric. However, volatility cascades are strongly asymmetric when moving from short to long horizons.

2.3 The Impact of Cyberattacks on Cryptocurrency Returns

There is a large body of literature which investigates different types of cyberattacks on blockchain technologies and the impact on the affected cryptocurrencies. For instance, Ramos et al. (2021) surveys the most common types of cyberattacks on Proof of Work (PoW) cryptocurrencies. Their empirical investigation covers three types of attack: 51 percent attack, where a group of miners gain control of more than 50 percent of the network's mining hash rate, hard forks, when a blockchain splits into two separate branches, and cryptocurrency wallet attacks.

Civitarese and Mendes (2018) conduct in event study to investigate the semi-strong form efficiency in the cryptocurrency market. They examine abnormal returns related to six negative events, classified as the so-called Fear, Uncertainty and Doubt (FUD) incidents. The chosen dates are related to critical technological failures and problems which received substantial media coverage. They show that cryptocurrencies quickly adjust to negative news announcements, possibly exhibiting a semi-strong form of market efficiency. In a related study Hasanova et al. (2019) surveys a broad range of blockchain technology cybersecurity vulnerabilities and provides a discussion of potential security countermeasures.

In contrast to the above mentioned studies which primarily investigate various vulnerabilities in blockchain architectures, a smaller subset of the literature places the focus on the theft of cryptocurrencies and cybersecurity breaches of cryptocurrency exchanges that are more relevant to the topic of this thesis.

Brown and Douglass (2020) investigate the effect that news of a cryptocurrency theft has on the price of cryptocurrencies. The thefts they consider include cryptocurrency thefts from digital exchanges as well as from large holders of cryptocurrencies. They apply event study methodology using a 3-days window from one day before to one day after each incident. Their study is limited to 16 events over the 2014 – 2019 period and 10 major cryptocurrencies. Surprisingly the reported results indicate that news of cryptocurrency thefts increases cryptocurrency prices of the affected cryptocurrency. The authors note that this finding is largely counter-intuitive as one would expect that news of cryptocurrency theft would decrease the value of cryptocurrencies because it shows the vulnerability of cryptocurrency storage methods.

Hu et al. (2020) analyse Bitcoin price dynamics before and after 30 cyberattacks over the 2012 – 2018 period that resulted in Bitcoin theft. The authors use USD amount stolen and proportion of the market volume to determine the size of the investigated hacks. They find a positive correlation between the volume of cryptocurrencies stolen as a proportion of the trading volume and the price two and three days before the hack. Moreover, 26 out of 30 investigated

incidents are associated with a price drop from the day before the incident to the day of the hack.

Lyócsa et al. (2020) investigate a number of related issues regarding the effect of news announcements on the volatility of Bitcoin. In particular they study whether news and sentiment about Bitcoin regulation, the hacking of Bitcoin exchanges and scheduled macroeconomic news announcements affect the realized variance of Bitcoin. In regards to cybersecurity related data they include 55 cryptocurrency-related cyberattacks, including the hacking on exchanges, online wallet providers and other types of cryptocurrency hacking. They report that the volatility of Bitcoin reacts most strongly to news on Bitcoin regulation, positive investor sentiment regarding Bitcoin regulation extracted using Google searches, and hacking attacks on cryptocurrency exchanges. Quantile regression results reveal that hacking attacks have particularly strong impact on the upper conditional distribution of Bitcoin volatility. On the other hand the volatility of Bitcoin is not influenced by most scheduled US macroeconomic news announcements, such as government budget deficits, inflation, or even monetary policy announcements.

Corbet et al. (2020a) examine the effect of cryptocurrency related security breaches using 60-min frequency data for the top eight cryptocurrencies and 17 largest cryptocurrency hacking events. The sample time period studied is September 2017 – August 2018. The return equation is modeled as a function of past returns, a number of pricing factors such as gold, VIX, oil, US equities and exchange rates, as well as cybersecurity breach dates via dummy (indicator) variables. This formulation is augmented with a dynamic conditional correlation (DCC) model to account for the time-varying nature of correlations and volatility. Hacking events are found to increase both the price volatility of the targeted cryptocurrency and broad cross-cryptocurrency correlations. Further, cybercrime events significantly reduce price discovery sourced within the hacked currency relative to other cryptocurrencies. Finally, abnormal returns associated with the hacks range between -2 percent to -24 percent, depending on the specific event. The abnormal returns are observed 4 hours before the actual hacking event and revert back to zero at the time and announcement of the hack.

Gandal et al. (2018) provide an interesting study of suspicious bot trading activity on Mt. Gox cryptocurrency exchange in which approximately 600,000 Bitcoins (BTC) were fraudulently acquired. It was later revealed that Mt. Gox exchange itself operated the fraudulent accounts in order to boost trading volume, collect extra trading fees, and cover up a 650,000 Bitcoin loss to hackers, which happened prior to the start of the reported bot trading. This internally orchestrated fraudulent trading was associated with a BTC price rise when the suspicious trades took place, compared to a slight decline on days without suspicious activity due to artificially boosted demand via bot trading. Gandal et al. (2018) hypothesise that the fraudulent bot trading acted as a signal to the market and encouraged others to enter and purchase Bitcoins.

Finally, Caporale et al. (2021) study dynamic linkages (interdependence) between cryptocurrencies and whether shifts in their spillover parameters (contagion) are associated with the occurrence of cyberattacks (contagion) over the period from August 2015 to January 2020. In particular, they study mean and volatility spillovers between three cryptocurrencies, including Bitcoin, ethereum and litecoin, and test for contagion effects, i.e. strengthening of linkages, during episodes of cyberattacks. Their database of cybercrime is catalogued by type, including

cybercrime, cyberespionage, cyberwarfare, and hacktivism, as well as by target such as crypto, government, industry and financial. When cyberattacks are not taken into account there is little evidence of spillovers between the returns of the investigated digital currencies, but there is strong evidence of volatility spillovers. In contrast, there is a downward shift (negative contagion) in the parameter measuring mean spillovers on the days associated with cyberattacks. As for linkages between the second moments, there are significant volatility spillovers from Bitcoin to Litecoin and Ethereum, whose size is magnified by cyberattacks.

3 Data Description

In the first part of the study, I compile a dataset on cryptocurrency exchange attributes in order to use them in predicting which digital markets will stay in business and which ones will close down. Some key predictor variables are cybersecurity features implemented by the exchanges, as well as a record of previous cyberattacks on the exchanges. In the second section I describe the data used to investigate the impact cyberattacks against cryptocurrency exchanges on the price/return dynamics of Bitcoin.

3.1 Predicting Cryptocurrency Exchange Closures

This dataset consists of cross-sectional data on 279 exchanges collected for the June 2010 – February 2022 time period. I construct the database by collecting information from publicly available sources, as well as incorporating some of the data published in previous studies such as Moore et al. (2018) and Oosthoek and Doerr (2020). In particular, I obtain information on security breaches from online lists compiled by Hackernews (2019), Selfkey (2019) and Slowmist (2021), and from other various media sources. Additional data on transaction volumes and exchange lifetimes is collected from online information portals and news websites such as coinmarketcap.com, coingecko.com, cryptowisser.com and coinpaprika.com. Lastly, each exchange website is manually inspected for information on cyber-security programs and any additional relevant information. To view the websites of closed exchanges, I rely on the Wayback Machine, which is described as a digital archive of the World Wide Web (archive.org).

The variable I aim to predict, i.e. target variable, is named *active* and is a binary variable signifying if an exchange remains active or has closed down, as defined below:

$$active_i = \begin{cases} 1 & \text{if cryptocurrency exchange } i \text{ remains active} \\ 0 & \text{if cryptocurrency exchange } i \text{ shuts down.} \end{cases} \tag{1}$$

The list of predictors comprises eight features including, i) *volume* – average daily traded volume in USD for each exchange, ii) *lifetime* – exchange lifetime in days and iii) *breach* – a binary variable which records if there has been a previous security breach or not for each exchange in the dataset. Amongst other predictors are binary variables representing whether or not each of the following four cyber-security measures is implemented iv) *two-factor* authentication, v)

bug-bounty program[3], vi) *security-audit*, and vii) *cold-storage*[4]. Nevertheless, not all exchanges provide information regarding all four security programs on their website. In such cases of missing data, and for the purpose of maximising the sample size, I take a conservative approach and code missing samples as 0, implying that the exchange for which the data is missing does not implement the security measure in question. This is a reasonable assumption, given that cryptocurrency investors worry about cyber risks and that digital exchanges compete on the basis of implemented security features. Finally, the dataset is completed with a variable capturing the extent of financial regulation in the country of origin of each exchange. Thus, the remaining predictor is viii) *aml/cft* – the anti-money laundering and combating the financing of terrorism index of Verdugo Yepes (2011) that measures the extent of a country's compliance with the anti-money laundering and combating the financing of terrorism (AML/CFT) international standard. Where an exchange operates in multiple countries, I take a conservative approach and classify it as operating in the country with the lowest *aml/cft* score.

Table 1 provides some summary statistics for the dataset.

Table 1: Descriptive Statistics

	mean	std	min	25%	50%	75%	max
active	0.52	0.50	0.00	0.00	1.00	1.00	1.00
breached	0.27	0.45	0.00	0.00	0.00	1.00	1.00
two-factor	0.90	0.31	0.00	1.00	1.00	1.00	1.00
bug-bounty	0.30	0.46	0.00	0.00	0.00	1.00	1.00
security-audit	0.28	0.45	0.00	0.00	0.00	1.00	1.00
cold-storage	0.78	0.41	0.00	1.00	1.00	1.00	1.00
aml/cft	27.32	6.72	11.90	23.33	28.33	33.67	35.33
volume	273.04	810.75	0.00	0.05	8.92	132.75	7344.85
lifetime	1474.57	922.24	19.00	794.00	1326.00	1945.50	3885.00

Notes: Dataset comprises 279 cryptocurrency exchanges; volume is measured in millions of USD.

As indicated by the first column of the table, about 52 percent of the 279 cryptocurrency exchanges contained in the database remain active, while 27 percent of the exchanges have suffered some form of a security breach. A majority of the exchanges implement two-factor authentication (90 percent of all exchanges) as well as cold storage facilities (78 percent). On the other hand, bug bounty and security audits are less commonly implemented by digital exchanges, with respective frequencies of 30 and 28 percent. The anti-money laundering and combating of financing of terrorism (*aml/cft*) index varies substantially and exhibits a mean value of 27.32 out of 49. I can also deduce the amount of skewness by comparing the difference

[3]Bug bounty is a program offered by websites and software developers by which individuals can receive recognition and compensation for reporting bugs, especially those pertaining to security vulnerabilities.

[4]Cold storage (cold storage wallet) is a hardware device used to store cryptocurrency that is kept offline, thus protecting the funds from unauthorized access, cyberattacks and other vulnerabilities to which a system that is connected to the internet is susceptible.

between the 25% and the 50% percentiles on the one hand, and the 50% and the 75% on the other hand. For instance, it is clear that the volume variable has a very long right tail.

Of particular interest are *lifetime* and *volume* variables which convey information about the relative success of the studied exchanges. The average lifetime for the sample of cryptocurrency exchanges appears to be about 1474.57 days, with a minimum of 19 days and a maximum of 3885 days. Thus, some exchanges have been exceptionally short lived. In addition, the standard deviation of *lifetime* is 922.24 days which is large relative to its mean value. The mean daily *volume* is USD 273.04 million and also varies substantially from USD 67.21 (displayed as 0.00 in USD millions) to USD 7,344.85 million. Overall the dataset appears to be highly heterogeneous in terms of exchange properties.

Next, I present pairwise correlations in Table 2 which contribute to prediction accuracy discussed in Section 5.

Table 2: Correlation Matrix

	active	breached	two-factor	bug-bounty	security-audit	cold-storage	aml/cft	volume	lifetime
active	1.00	-0.09	0.33	0.38	0.44	0.39	-0.09	0.28	0.53
breached	-0.09	1.00	-0.21	0.04	0.09	-0.16	0.01	0.01	0.03
two-factor	0.33	-0.21	1.00	0.22	0.19	0.53	-0.07	0.11	0.35
bug-bounty	0.38	0.04	0.22	1.00	0.36	0.21	0.04	0.17	0.14
security-audit	0.44	0.09	0.19	0.36	1.00	0.25	0.01	0.17	0.32
cold-storage	0.39	-0.16	0.53	0.21	0.25	1.00	-0.07	0.15	0.26
aml/cft	-0.09	0.01	-0.07	0.04	0.01	-0.07	1.00	-0.04	-0.01
volume	0.28	0.01	0.11	0.17	0.17	0.15	-0.04	1.00	0.08
lifetime	0.53	0.03	0.35	0.14	0.32	0.26	-0.01	0.08	1.00

Notes: Computations based on a dataset comprising 279 cryptocurrency exchanges.

Considering the correlations between the target variable and various predictors provided in the first row of the table, I observe that the highest correlation of 0.53 is recorded between the target *active* and the predictor *lifetime*. Other features such as *security-audit* and *cold-storage* also exhibit relatively large and positive correlations with the target, which are respectively 0.44 and 0.39. These are followed in magnitude by *bug-bounty* and *two-factor* with the correlations of 0.38 and 0.33. Additionally, *volume* also has a positive and moderate correlation with *active* of 0.28, while *breached* and *aml/cft* variables are negatively correlated with the target variable. Thus, it would appear that implementing cyber-security features, having a longer trading track record and a greater transaction volume is positively associated with exchanges that succeed at remaining active. On the other hand, experiencing a security breach and operating in countries with greater emphasis on anti-money laundering regulation is negatively related with the variable *active*, although these correlations are rather small in magnitude. In the second row of Table 2, I observe that experiencing a security breach is negatively related to *two-factor* and *cold-storage*, as expected. These correlations are, however, not large in magnitude with the estimates of -0.21 and -0.16.

Rows 3 – 6 suggest that the four cybersecurity measures are positively correlated, implying that digital exchanges that implement sound security practices tend to do so across multiple measures. The lowest of these correlations (0.19) is found between *security-audit* and *two-factor*, which can be explained by the fact that two-factor authentication is easy to implement internally while a security audit is costly and requires engagement with external security auditors. In contrast, the highest correlation of 0.53 is recorded between *cold-storage* and *two-factor* features. Although the magnitude of this correlation may present some difficulty in disentangling the individual effects of *cold-storage* and *two-factor* on the target *active*, it will have no impact on the overall classification performance[5]. Given that I aim to maximize the forecasting ability, I decide to leave all four security features in the dataset. Lastly, while *volume* exhibits relatively low correlations with other predictors, *lifetime* seems to be moderately and positively correlated with *two-factor* and *security-audit*.

3.2 Evaluating the Impact of Cyberattacks Against Cryptocurrency Exchanges on Bitcoin Returns

In this subsection I define the dataset that is used in the second part of this study to assess the impact of cyberattacks against cryptocurrency exchange on Bitcoin returns.

I collect daily time series for the following five variables: i) dates on cybersecurity breaches of major cryptocurrency exchanges, ii) the price of Bitcoin (BTC)[6], iii) a US Dollar index[7], iv) the price of gold[8], and v) a value-weighted US equities index[9]. Variables ii) - v) are expressed as percent returns $r_t = 100 \times (P_t/P_{t-1} - 1)$, which are used in equations (6) and (7) in Section 4.

The main variable – the dates of cryptocurrency exchange breaches – is related to the variable *breached* discussed in the previous section and now collects breach dates as a time series across all exchanges included in our dataset (which occur over the 2012 – 2021 period). It is constructed as an indicator variable that takes the value one on the dates of the recorded cyberbreaches and the value zero for all other dates as follows:

$$I_t = \begin{cases} 1 & \text{if } t \text{ is a breach date of a cryptocurrency exchange,} \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

In total there are 75 cybersecurity breaches which are accounted for by I_t above. Note that this is also the number of the exchanges for which I have records of in Table 1 above, i.e. 0.27×279.

[5]Finding high correlations between features is termed multicollinearity and reduces the precision of the estimated coefficients. Nevertheless, multicollinearity does not impact the precision of the estimated linear combination of the features which is used to generate predictions.

[6]The price of Bitcoin is as recorded on Bitstamp, one of the longest running cryptocurrency exchanges that was founded in 2011. Source: https://bitcoincharts.com/charts/bitstampUSD#a1gWMAzm1g5zm2g10zl.

[7]Nominal broad US dollar index. Source: Federal Reserve Bank of St. Louis https://fred.stlouisfed.org/series/DTWEXBGS.

[8]COMEX gold. Source: https://finance.yahoo.com/quote/GC%3DF/history?p=GC%3DF.

[9]This is a value-weight return index of all CRSP firms incorporated in the US and listed on the NYSE, AMEX, or NASDAQ as provided by Kenneth R. French from https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

The total amount stolen across these 75 major incidents is USD $2.25 billion. This estimate is nevertheless a conservative figure as not all recorded cybersecurity breaches have disclosed stolen amounts. The final dataset covers the time period from January 3, 2012 to December 28, 2021, and contains 2470 daily observations. As explained in the previous section this data is obtained from publicly available sources such as the lists compiled by Hackernews (2019), Selfkey (2020) and Slowmist (2022). The veracity of these preliminary dates is checked manually against various media sources and digital exchange websites.

The fact that traditional markets do not trade on weekends and public holidays, while Bitcoin does, implies that in order to combine Bitcoin returns with the control variables, such as the US equities index, I had to remove weekend and public holiday observations from the dataset. While reducing the sample size this is an important step because omitting relevant control variables can significantly bias estimation results. Similar pricing factors have been found to be important for the formation of Bitcoin price in a number of studies such as Corbet et al. (2020b), Corbet et al. (2020a), Bouri et al. (2018), Klein et al. (2018), etc.

Table 3 provides some basic summary statistics for the dataset. As illustrated in the first column of the table Bitcoin experienced the highest daily mean return of about 0.46 percent over the time period. This is accompanied by equally high volatility of 4.90 as measured by the standard deviation. The exceptionally high levels of risk are also evident from the minimum and maximum daily return values of -48.52 and 40.14 percent, respectively. US equities have rank second in terms of both the return and risk, but when compared to Gold they provide a better risk-reward ratio with marginally higher risk and seven fold increase in return. US Dollar index and Gold exhibit more moderate returns of 0.01 percent while Gold is three times as risky with 0.99 standard deviation. Finally, Breach Day Indicator binary variable has a mean value of 0.03 signifying that about 3 percent of the sample observations are associated with cybersecurity breaches.

Table 3: Descriptive Statistics

	mean	std	min	25%	50%	75%	max
Bitcoin	0.46	4.90	-48.52	-1.37	0.26	2.25	40.14
US Dollar (TWI)	0.01	0.31	-2.07	-0.16	-0.00	0.18	1.94
Gold	0.01	0.99	-9.35	-0.45	0.01	0.51	5.95
US Equities	0.07	1.05	-11.99	-0.34	0.09	0.55	9.35
Breach Day Indicator	0.03	0.17	0.00	0.00	0.00	0.00	1.00

Notes: Statistics computed from daily data over the January 3, 2012 – December 28 time period.

Figure 1 depicts the price and returns evolution of Bitcoin over time, as well as the dates of cybersecurity breaches. As illustrated, the Bitcoin price has experienced incredible growth over the sample period starting at $5.29 in January 2012 and finishing at $47,543.30 in December 2021. Over the sample period, I also observe cybersecurity breaches experienced by crypto exchanges, as depicted by vertical grey lines in the same figure. 2014 to mid 2015 and 2018 to

2021 periods appear to experience especially high frequencies of cyberattacks. While there are some coincidences of large negative Bitcoin returns and breach days, it is difficult to judge their relationship based on the information provided in Figure 1 alone, and a formal statistical test is required to draw any firm conclusions.

Figure 1: Bitcoin Prices/Returns and Cyberbreaches of Cryptocurrency Exchanges

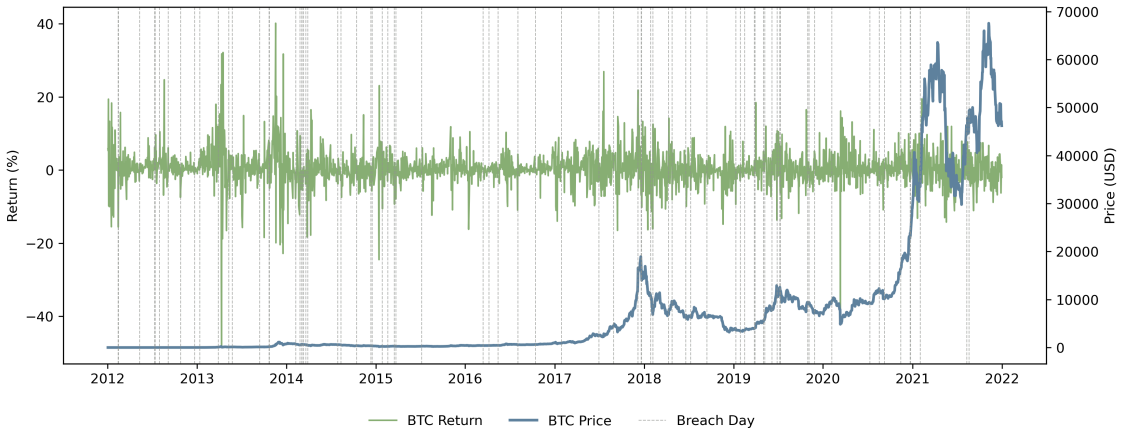

4 Methodology

This section is divided into two parts. First, I discuss an empirical method for predicting which cryptocurrency exchanges will remain open and which ones will close based on various exchange attributes. In the second subsection, I present the methodology used to evaluate the impact of cyberattacks against cryptocurrency exchanges on the return of Bitcoin.

4.1 Predicting Cryptocurrency Exchange Closures

The empirical method consists of three steps: i) training and optimising ML algorithms, ii) evaluating in-sample (training dataset) and out-of-sample (test dataset) classification performance and ranking the algorithms according to their predictive ability, and iii) examining feature importance and determining which predictors contribute to forecasting ability. I start by discussing the problem of predicting which digital exchanges will remain active and which ones will face closure.

4.1.1 The Prediction Problem

This thesis aims to predict which cryptocurrency exchanges will remain active and which will go out of business, conditional on a set of relevant predictor variables. This task may be formulated as a classification problem where the target variable is defined in (1).

Forecasts of the target variable *active* (y_i) are denoted as $\hat{y}_i^a$ and are generated based on eight available predictor variables ($x_{1i}, x_{2i}, \ldots, x_{8i}$) which are discussed in detail in Section 3.1. Thus, the forecasts are constructed according to the following equation

$$\hat{y}_i^a = \phi^a(x_{1i}, x_{2i}, \ldots, x_{8i}), \tag{3}$$

where ϕ^a is a function describing the relationship between the forecast and predictor variables that depends on which forecasting algorithm (denoted by a) is used.

While there are a plethora of classifiers one could apply to the prediction problem, in this investigation, I decide to compare the performance of ten popular ML algorithms. These are as listed below:

1. Logistic Regression;
2. Decision Tree;
3. Random Forest;
4. Support Vector Machine;
5. Multi-layer Perceptron;
6. KNeighbors;
7. Naive Bayes;
8. AdaBoost;
9. ExtraTrees;
10. Equally-weighted ensemble of the above 9 classifiers.

These models are flexible and capable of capturing complicated relationships between the target variable and relevant features. A brief description of each classifier is provided next.

Logistic regression is one of the earliest and most widely employed methods used for modelling of binary dependent variables, see, e.g. Wilson and Worcester (1943). It specifies the conditional probability of success given the vector of predictors x_i, as a sigmoid function of the following form $P(y_i = 1|x_i) = \frac{1}{1+e^{-w'x_i}}$, where w refers to the vector of weights, including the intercept. Prediction of class membership is then generated as follows

$$\hat{y}_i^{LR} = \begin{cases} 1 & \text{if } P(y_i = 1|x_{1i}, x_{2i}, \ldots, x_{8i}) \geq 0.5 \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

Since I do not implement any regularization in the logistic regression, I am also able to estimate standard errors, and thus gauge the statistical significance of the estimated coefficients.

A *Decision tree* classifier has the ability to construct complex decision boundaries by dividing the feature space into rectangles. The decision tree consists of a root node, decision nodes and terminal nodes, as illustrated in Figure 2. These nodes are formed by starting at the tree root and splitting the data on the feature that results in the largest information gain (IG). The splitting procedure is repeated at each decision node until the leaves either contain elements from only

one class or by setting a limit for the maximal depth of the tree (which avoids overfitting). In the application I employ grid search cross-validation to optimise two hyperparameters i) tree depth, and ii) criterion used to compute IG.

Figure 2: A Decision Tree Classifier

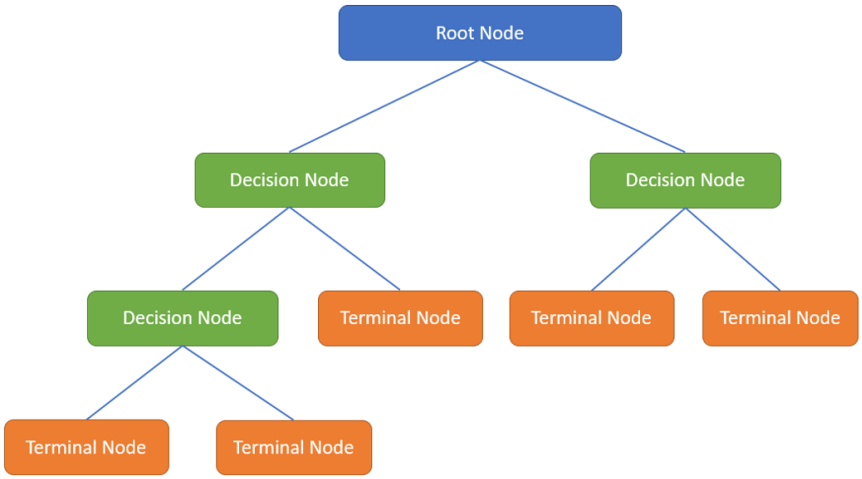

A *random forest* classifier is an ensemble of decision trees. Random forests combine predictions of multiple decision trees by averaging across their estimated probabilities, thus reducing the degree of overfitting. A random forest of k trees may be constructed via the following algorithm:

i) Draw a random bootstrap sample of size n from the training dataset (with replacement);

ii) Grow a decision tree from the bootstrap sample:

   (a) Randomly select d features without replacement;

   (b) Build a tree using these d features;

iii) Repeat i) - ii) k times;

iv) Combine classifiers by averaging their probabilistic predictions. Assign class label according to greatest probability.

In the application, I optimize three random forest hyperparameters: i) maximum tree depth, ii) number of trees k, and iii) criterion used to compute IG.

A *Support vector machine* classifier is an algorithm designed to be robust to outliers. It works by maximizing the margin, which is determined by the distance between the decision boundary and the training examples that are closest to the boundary, i.e. support vectors, as illustrated in Figure 3. I implement the algorithms with L2 regularization and optimize the regularization strength parameter. In addition, I cross validate the kernel function as a hyperparameter across

the following values {linear, rbf, polynomial, sigmoid}. This allows the support vector machine to accurately classify linearly inseparable (nonlinear) datasets.

Figure 3: A Support Vector Machine Classifier

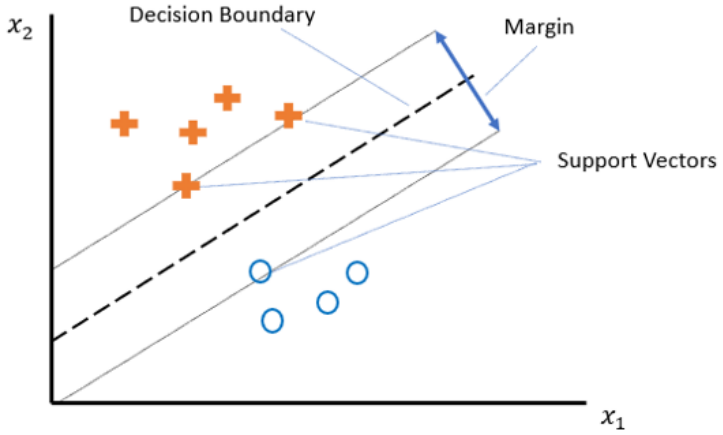

A *multilayer perceptron* (MLP) is a feed-forward neural network, consisting of three or more layers: input layer, hidden layer(s) and output layer, see e.g. Azar and El-Said (2013), as depicted in Figure 4. The leftmost layer, known as the input layer, consists of a set of neurons representing the input features. Each neuron in the hidden layer transforms the values from the previous layer with a weighted linear summation $w_1x_1 + w_2x_2 + \cdots + w_nx_n$ followed by a non-linear activation function $g(.) : R \to R$, for instance the hyperbolic tan function. The output layer receives the values from the last hidden layer and transforms them into output values. Thus the units in the hidden layer are fully connected to the input layer, and the output layer is fully connected to the hidden layer. If such a network has more than one hidden unit, it is called a deep artificial neural network (NN). In the application, I employ 10-fold cross validation to find the optimal values for the number of hidden layers, choosing between 1 and 2 layers. The number of neurons in the hidden layers is set to 8, the number of features in the data set. In addition, I also optimise the strength of L2 regularisation and the activation function, which is selected from the following set {identity, logistic, tanh, and relu}.

Figure 4: A Hypothetical Multilayer Perceptron

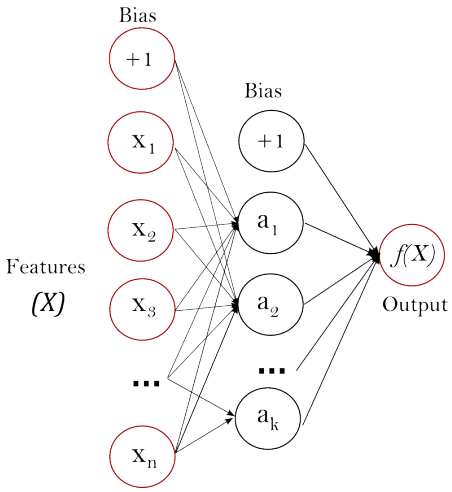

The next method I discuss is *KNeighbors* (nearest neighbors) that finds a predefined number of training samples closest in the distance to the new point and predicts the label from these. The algorithm may be described as follows:

- Choose the number of neighbors (k) and a distance metric;
- Find k nearest neighbors of the data example I need to classify;
- Assign the class label by majority vote.

Naive Bayes is an algorithm that applies Bayes' theorem with the assumption of conditional independence between each pair of features given the value of the class variable,
i.e. $P(x_i|y, x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n) = P(x_i|y)$. Substituting this assumption into the Bayes' theorem

$$P(y|x_1, \ldots, x_n) = \frac{P(y)P(x_1, \ldots, x_n|y)}{P(x_1, \ldots, x_n)} \tag{5}$$

I obtain

$$P(y|x_1, \ldots, x_n) = \frac{P(y)\prod_{i=1}^{n} P(x_i|y)}{P(x_1, \ldots, x_n)}.$$

Since $P(x_1, \ldots, x_n)$ is constant the classification rule may be expressed as follows:

$$\hat{y} = \underset{y}{\operatorname{argmax}} P(y)\prod_{i=1}^{n} P(x_i|y)$$

where $P(y)$ and $P(x_i|y)$ can be estimated using Maximum A Posteriori (MAP) method. In this application, I use Gaussian distribution to model $P(x_i|y)$ and optimize the variance smoothing

parameter, which is the portion of the largest variance of all features that are added to variances for calculation stability.

Boosting methods focus on observations that are difficult to classify. An *AdaBoost* classifier is a meta-estimator that begins by fitting a classifier on the original dataset and then fits additional copies of the classifier on the same dataset, but adjusts the weights of incorrectly classified instances such that subsequent classifiers focus more on complex cases. The base classifiers I use in the AdaBoost are decision tree stumps, i.e. one-level decision trees where the split at the root level is based on a specific attribute/value pair. I optimise the learning rate and the number of base estimators via cross validation.

An *ExtraTrees* (extremely randomised trees) classifier is similar to the random forest, but as its name suggests, the degree of randomness employed in splitting the data is increased. As in random forests, a random subset of candidate features is used, but instead of looking for the most discriminative thresholds, thresholds are drawn at random for each candidate feature, and the best of these randomly-generated thresholds is picked as the splitting rule. I optimise the number of base estimators used in *ExtraTrees*.

The final model applied is a majority voting *ensemble* which joins the nine previously discussed classifiers. It combines the predictions from all the other models by predicting the label that has been predicted by the majority of classifiers (received at least 50% of votes).

In order to improve the convergence properties of the ML algorithms, I normalize all continuous predictors to have zero mean and unit variance, while the binary features are left unchanged. All algorithms are implemented in Python using scikit-learn libraries.

4.1.2 Cross-Validation and Performance Evaluation

In order to assess forecasting performance on an independent dataset, I divide the data into training and test subsamples. The training dataset contains 70 percent of the data (195 observations), while the test dataset consists of the remaining 30 percent (84 observations). When splitting the data, I preserve the proportions of examples in each class by stratifying the data according to the target variable. Alternative 80:20 and 60:40 splits between the training and test datasets have been tried and result in similar classification performance.

The models are first trained and their hyperparameters optimized using the training dataset and K-fold cross-validation where $K = 10$. In the second step, I compute both in-sample (training dataset) and out-of-sample (test dataset) forecasting performance according to the following measures: i) classification accuracy, ii) precision, iii) recall, and iv) F1 score. These metrics gauge somewhat different aspects of forecasting ability that cryptocurrency investors may care about. Denoting true positives as TP, true negatives as TN, false positives as FP and false negatives as FN, I explain the four performance metrics as follows.

i) Accuracy $= \frac{\text{number of correctly classified examples}}{\text{sample size}} = \frac{TP+TN}{TP+TN+FP+FN}$. Classification accuracy is defined as the ratio of correctly predicted examples to the total sample size and is probably the most commonly used measure of classification performance. Nevertheless, it has a disadvantage that in situations where there is a class imbalance, the model can predict the value of the majority class for all samples and still achieve a high classification accuracy.

ii) Precision $= \frac{TP}{TP+FP}$. Precision computes the ratio of true positives to all positively labelled (predicted) examples. It answers the question of how many exchanges survive out of all the exchanges which are predicted to survive.

iii) Recall $= \frac{TP}{TP+FN}$. Recall is the ratio of correctly predicted positive example to all positive samples. It tells us how many exchange are predicted to survive out of all exchanges that truly survive.

iv) F1 Score $= 2 \times \frac{\text{Recall} \times \text{Precision}}{\text{Recall} + \text{Precision}}$. F1 score is computed as the weighted average of precision and recall and aims to balance these two metrics.

Using the above four measures I compute in-sample (training dataset) and out-of-sample (test dataset) classification ability of each of the ten models discussed previously. The algorithms are then ranked according to their performances.

Having discussed the methodology employed in predicting which cryptocurrency exchanges will remain active and which ones will close down I now turn to explain the method used to accomplish the second task of this thesis, namely assessing the impact of cyberattacks against cryptocurrency exchanges on Bitcoin returns.

4.2 Evaluating the Impact of Cyberattacks Against Cryptocurrency Exchanges on Bitcoin Returns

In order to investigate the impact of digital exchange cyberattacks variable I_t, defined in equation (2), on Bitcoin returns r_t^{BTC} I consider the following regression:

$$r_t^{BTC} = c + \delta_0 I_t + \sum_{k=-3}^{-1} \delta_k I_{t+k} + \sum_{k=1}^{3} \delta_k I_{t+k} + \phi r_{t-1}^{BTC} + \varepsilon_t. \tag{6}$$

The contemporaneous effect of cyberbreaches is captured by the parameter δ_0, while I control for three days before and three days after the incidents using the leads ($k = 1, 2, 3$) and lags ($k = -1, -2, -3$) of I_t. Finally, I include a lagged value of BTC return itself in order to account for possible autocorrelation in the regression equation. This type of model is widely applied in event-driven analysis across many applications, such as Nikkinen and Sahlström (2004) and Chen and Clements (2007), and more recently in Pyo and Lee (2020) in the context of Bitcoin.

While Bitcoin is referred to as a cryptocurrency, it is also often considered to be an asset, e.g. Glaser et al. (2014). As such, its returns may be driven by factors similar to those found in the asset pricing literature. Omitting relevant factors from (6) may bias the estimate of δ_0 and make results unreliable. Indeed, studies such as Corbet et al. (2020b), Corbet et al. (2020a), Bouri et al. (2018), Klein et al. (2018), etc., account for pricing factors in the context of Bitcoin returns. Thus, I augment the model in (6) using the returns on a US equities, a US Dollar index, and gold. These factors control for the US equities market, exchange rates and commodities, respectively, and are amongst the factors used recently in Borri et al. (2022) and Corbet et al. (2020a). The extended model then becomes

$$r_t^{BTC} = c + \delta_0 I_t + \sum_{k=-3}^{-1} \delta_k I_{t+k} + \sum_{k=1}^{3} \delta_k I_{t+k} + \theta_1 r_t^{US\ Equities} + \theta_2 r_t^{US\ Dollar} + \theta_3 r_t^{Gold} + \phi r_{t-1}^{BTC} + \varepsilon_t. \quad (7)$$

I complete the specification by assuming zero conditional mean for the error term ε_t, i.e. $\mathbb{E}(\varepsilon_t|\varepsilon_{t-1}, \varepsilon_{t-2}, ...) = 0$, and considering three alternative assumptions regarding the conditional variance $\text{var}(\varepsilon_t|\varepsilon_{t-1}, \varepsilon_{t-2}, ...) = \sigma_t^2$

1. $\sigma_t^2 = \omega$ for all t (OLS regression model);
2. $\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2$ (GARCH model);
3. $\sigma_t^2 = \omega + (\alpha + \gamma I_{[\varepsilon_{t-1}<0]})\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2$ (GJR model).

Model 1 above makes the simplest assumption of constant conditional variance as is commonly assumed in the ordinary least squares (OLS) regression model. However, existing literature reports that Bitcoin returns exhibit time-varying conditional variances, see, e.g., Dyhrberg (2016) and Katsiampa (2017), and this is what the remaining two specifications attempt to capture. Model 2 is a standard GARCH(1,1) model of Bollerslev (1986), where the current value of the time-varying variance σ_t^2 is a function of past variance and a squared past error term, while model 3 is the GRJ(1,1,1) specification of Glosten et al. (1993) that accounts for asymmetries associated with negative returns via its γ parameter. Accounting for the time-varying volatility property is important because it improves estimator efficiency and consequently leads to more powerful statistical tests.

5 Results and Discussion

The results section is divided into two main parts. In the first subsection, I present the results of modeling and predicting which cryptocurrency exchanges will remain active and which markets will close down. In the second part I provide my findings of investigating the impact of cyberattacks against cryptocurrency exchanges on Bitcoin price/return.

5.1 Predicting Cryptocurrency Exchange Closures

I start with a discussion of classification performance results, which are then followed by the analysis of feature importance and a visualisation of some predictions.

5.1.1 Measuring Classification Performance

Table 4 presents four measures of in-sample classification performance, which are computed using the training dataset. While the data is split according to the 70:30 percent ratio between the training and test datasets, alternative schemes result in similar classification performances.

First, I note that all ten algorithms achieve satisfactory performance according to the results presented in the table. The best performing algorithm is the ensemble classifier which combines the predictions from the other nine models and reaches in-sample accuracy of 0.959. In the second place is the random forest classifier, while the third place is shared by the multilayer perceptron and support vector classifier with the accuracy of 0.949. The remaining six algorithms reach accuracies ranging from 0.836 (ExtraTrees) to 0.944 (AdaBoost). Thus, the difference between the highest and the lowest classification accuracy is about 12.3 percent when the accuracy is computed using the in-sample (training) dataset.

Table 4: In-sample Forecasting Performance (Training Dataset)

Algorithm	Accuracy	Precision	Recall	F1 Score
Ensemble	0.959	0.960	0.960	0.960
Random Forest	0.954	0.960	0.950	0.955
Multilayer Perceptron	0.949	0.950	0.950	0.950
Support Vector	0.949	0.942	0.960	0.951
AdaBoost	0.944	0.941	0.950	0.946
KNeighbors	0.938	0.924	0.960	0.942
Decision Tree	0.928	0.958	0.901	0.929
Naive Bayes	0.872	0.913	0.832	0.870
Logistic Regression	0.867	0.879	0.861	0.870
ExtraTrees	0.836	0.822	0.871	0.846

Notes: Performance metrics are computed from the training dataset consisting of 195 samples (70 percent of the dataset).

Although Table 4 sorts values according to classification accuracy, all four performance measures are largely consistent in their rankings. For instance, according to the first row of Table 4, the ensemble classifier ranks first in terms of accuracy, recall and F1 score. When ranked by precision, the ensemble classifier shares the first place with random forest. Similarly, the distinction between the second (random forest) and the third (multilayer perceptron) place is evident in three out of four metrics, with recall being the only measure that ranks random forest and multilayer perceptron equally. Nevertheless, most of the classifiers record performance measures of roughly similar magnitudes.

Having explored in-sample classification performance, I now turn to out-of-sample metrics provided in Table 5, computed using the test dataset. These results provide a better representation of the true predictive ability since the test dataset has not been used for the purpose of fitting the algorithms or optimizing hyperparameter values. While there is some decrease in classification performance relative to the training dataset, the out-of-sample accuracy figures reported in Table 5 range between 0.738 and 0.857, suggesting that predictive models generalize well to unseen test data.

The best performing model listed in row one of Table 5 is the random forest classifier with the accuracy of 0.857. While this value signifies a high level of predictability, it also represents about a 10 percent decrease in the accuracy of the best performing model when compared to

Table 5: Out-of-sample Forecasting Performance (Test Dataset)

Algorithm	Accuracy	Precision	Recall	F1 Score
Random Forest	0.857	0.848	0.886	0.867
Multilayer Perceptron	0.833	0.857	0.818	0.837
Ensemble	0.833	0.857	0.818	0.818
Decision Tree	0.833	0.841	0.841	0.841
Logistic Regression	0.810	0.868	0.750	0.805
Naive Bayes	0.798	0.909	0.682	0.779
ExtraTrees	0.798	0.865	0.727	0.790
Support Vector	0.774	0.791	0.773	0.782
AdaBoost	0.750	0.683	0.977	0.804
KNeighbors	0.738	0.789	0.682	0.732

Notes: Metrics are computed on test dataset consisting of 84 samples (30 percent of all data)

the results in Table 4. This is however largely expected as most machine learning models tend to be somewhat overfitted on training data. The ensemble model, which previously recorded the best classification performance, now shares the second place with the multilayer perception according to accuracy, precision and recall. When the F1 score is computed using the precision and recall values, which are not rounded to three decimal places multilayer perceptron ranks second and above the ensemble model.

While the best three models in Table 4 remain in the top positions when evaluated on test data, although in permuted places, some algorithms drop in rank significantly. The worst performing model in out-of-sample comparison appears to be KNeighbors, with the recorded accuracy of 0.738. This value may be compared to the training set accuracy of 0.938 when KNeighbors ranked in sixth place, which suggests that this model is substantially overfitted. A similar finding holds for AdaBoost and the support vector classifier.

Given that random forest exhibits good performance both in-sample and out-of-sample according to multiple performance criteria, I examine its classification results in more detail by considering the confusion matrix provided in Figure 5.

Out of the total of 84 samples contained in the test dataset, 44 exchanges that remain active (class 1), and 40 exchanges have closed down (class 0). As can be seen from the second row, 5 of the 44 active exchanges are misclassified as facing closure by the random forest classifier. This corresponds to the recall value of 0.886 presented in Table 5. In contrast, of the 40 exchanges which went out of business I successfully predict 33, while 7 exchanges are misclassified as remaining active. This results in a true negative rate (specificity) of 0.825. Thus, while I am able to separate the classes with high accuracy, a certain amount of risk still remains when predicting which exchanges will close down and which ones will remain active.

In order to gain further insight into the problem, I consider which features contribute most to the reported classification ability.

Figure 5: Confusion Matrix – Random Forest Algorithm and Test Data

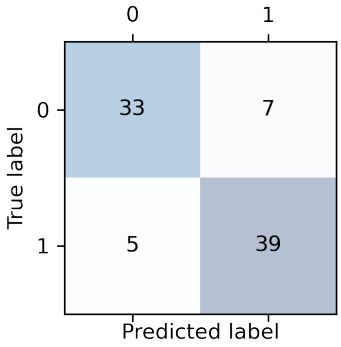

5.1.2 Feature Importance

Feature importance refers to the usefulness of predictors in forecasting the target variable. However, there is no single method to measuring feature importance that can be applied to all algorithms. For instance, the support vector classifier tackles potential nonlinearities via kernel methods which make it difficult to obtain even the simplest measures of feature importance such as the magnitude of the estimated weight coefficients[10]. Nevertheless, amongst the list of the classifiers which I implement here three models have well defined measures of feature importance that are easily computed. Should evidence from multiple algorithms suggest that a certain feature is "important", then I can have greater confidence in the impact of that predictor.

Figure 6: Feature Importance According to Random Forest Classifier

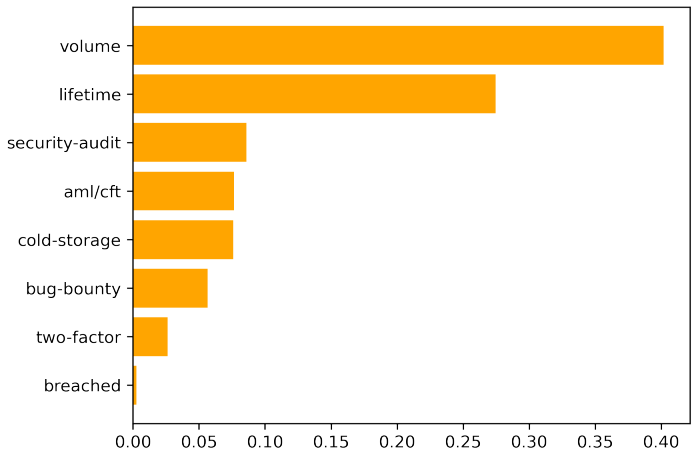

[10]Considering the magnitude of estimated weight parameters provides information about relative feature importance when the features are measured on the same scale (standardised in some way).

Figure 6 presents the estimates of Gini importance computed from the random forest classifier. These quantities are calculated as normalised reductions in node impurity (Gini impurity) resulting from every feature and then averaged across all estimated trees. As evident from the figure, *volume* (in USD) appears to be the main predictor used in separating which exchanges will remain active and which markets will close. The second most important feature is a *lifetime*. These two top predictors are followed by three cybersecurity features, namely, *security-audit*, *bug-bounty* and *cold-storage*, as well as by a measure of anti-money laundering regulation for country of origin, i.e. *aml/cft*. In contrast, *two-factor* has a much smaller effect, while *breached* seems to have a negligible impact on classification ability. The result regarding *two-factor* does not necessarily imply that two-factor authentication bears no importance for cybersecurity of digital markets but is likely to be an artifact of the sample composition, 90 percent of which implements two-factor authentication.

Next, I consider feature importance according to the decision tree classifier depicted in Figure 7. Here the importance of each feature is computed as the total reduction of Gini impurity resulting from that feature (similar to previously discussed random forest feature importance). I confirm the importance of *volume*, *lifetime*, *cold-storage* and *bug-bounty* predictors, which are also featured in Figure 6. However, in the case of the decision tree classifier the importance of *volume* relative to the other three predictors increases substantially. In fact, it is more than double the importance of *lifetime*, which is the second most important predictor. *Cold-storage* and *bug-bounty* play a small role in classifying exchange closures, while the remaining four predictors play no role in the decision tree classifier.

Figure 7: Feature Importance According to Decision Tree Classifier

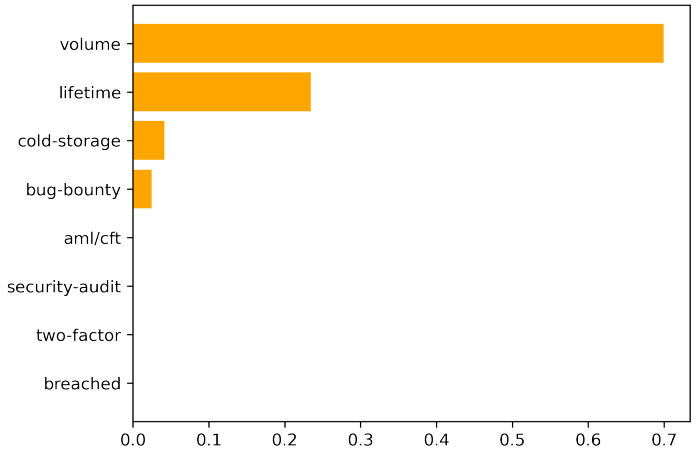

Lastly, I look at the logistic regression model. Table 6 reports marginal effects and their p-values, which provide a different perspective on feature importance to what I discussed above. Marginal effects measure how the predicted probability of a binary outcome changes with a change in a risk factor. For instance, I can look at how the probability of remaining active

changes with a 1-unit increase in (normalized) volume or exchange with a security audit versus an exchange without it. Using this approach, I can comment on both the magnitude of the impact for each predictor, i.e. the size of the marginal effect, as well as on their statistical significance.

Considering the p-values reported in the last column of the table, I see that *volume*, *security-audit*, *cold-storage*, *lifetime* and *bug-bounty* all exhibit statistically significant coefficients at the 5 percent level. All of these estimated effects are positive, implying that they increase the probability of remaining active. For instance, increasing the (normalised) *volume* feature by one unit will increase the probability of remaining active by 0.566, while a 1-unit increase in a (normalized) *lifetime* will result in a 0.147 change in the same probability. Of the binary variables, I see that implementing *cold-storage*, *bug-bounty*, and *security-audit*, respectively, result in 0.161, 0.145 and 0.172 improvements in the probability of remaining in business. These variables also played an important role in the random forest classifier, while the decision tree classifier identified a subset of them. *Breached* feature, which records the incidence of previous security breaches, is negative and statistically significant at the 5 percent level. This suggests that the digital markets which have previous experience with cyberattacks are more likely to close down and is consistent with the findings reported in Moore et al. (2018). Interestingly the results of the decision tree and random forest classifiers presented above do not corroborate this finding.

Table 6: Marginal Effects Estimated by Logistic Regression

	dy/dx	Std. Err.	z-stat.	p-value
volume	0.566	0.148	3.826	0.000
security-audit	0.172	0.045	3.859	0.000
cold-storage	0.161	0.081	1.988	0.047
lifetime	0.147	0.019	7.587	0.000
bug-bounty	0.145	0.042	3.433	0.001
aml/cft	-0.012	0.021	-0.538	0.591
breached	-0.109	0.054	-2.000	0.045
two-factor	0.750	14.376	0.052	0.958

Notes: The columns present i) marginal effects, ii) standard errors, iii) z-statistics and iv) p-values.

The remaining features, *two-factor* and *aml/cft* are not statistically significant at any conventional level of significance. However, as I can see *two-factor* variable has a large and positive estimated coefficient while *aml/cft* exhibits a small negative coefficient. The insignificance of *two-factor* is likely to be due to relatively high correlations between this variable and other security features making it difficult to disentangle individual effects (see Table 2), and the fact that about 90 percent of the digital markets in the sample implement two-factor authentication.

Lastly, in Figure 8, I plot the predicted and realised samples from the entire dataset of 279 exchanges against the three features designated as important by multiple algorithms – namely *volume*, *lifetime* and *cold-storage*. While the presented predictions are generated using

all eight features, the visualisation illustrates the relationship between the forecasts and the plotted predictors in a 3-dimensional subspace. As can be seen from the graph, the predictions (smaller solid circles) fall inside the realized data samples (larger transparent circles). More importantly, the colours of the larger and smaller circles mostly match, indicating a high degree of classification accuracy.

Figure 8: Visualizing Predictions in a 3-dimensional Subspace

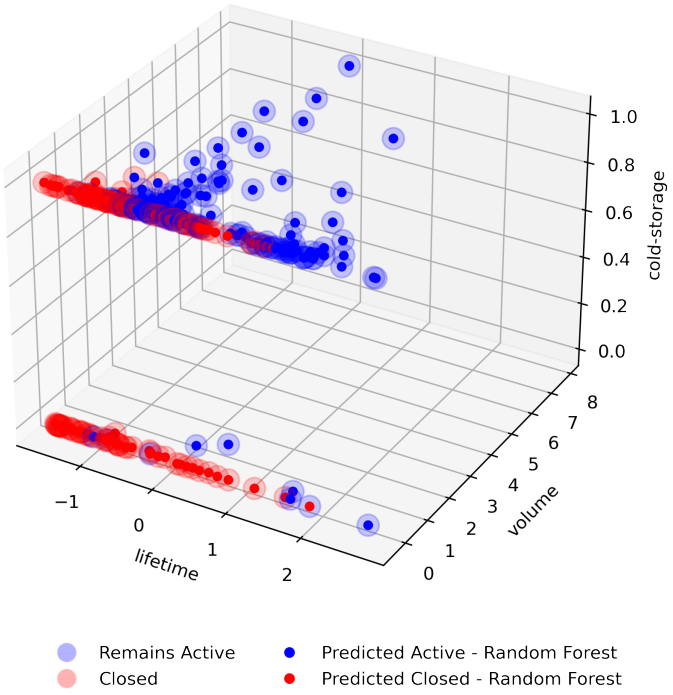

5.2 Evaluating the Impact of Cyberattacks Against Cryptocurrency Exchanges on Bitcoin Returns

In this section I present the empirical results of the models specified in (6) and (7) which aim to assess the magnitude and significance of the impact of cyberattacks on the returns of Bitcoin.

Table 7 provides the coefficients of the model specified in (6) and estimated in combination with the three alternative variance specifications discussed in the previous section. The parameter of primary interest is δ_0 which multiplies the contemporaneous indicator variable I_t capturing the incidences of cyberattacks against digital exchanges. Examining δ_0 across the three estimated models, I observe that it is consistently negative and statistically significant in all cases, at the 5% level. As expected, the effect varies across the models somewhat and ranges between -1.307% (when estimated by OLS) and -1.425% (in the case of the GJR model). The

interpretation of this result is that, at the 5% significance level, there is a statistically significant decrease in Bitcoin return on the dates corresponding with cyberattacks against cryptocurrency exchanges.

Table 7: Parameter Estimates

	OLS		GARCH		GJR	
	Coef.	p-value	Coef.	p-value	Coef.	p-value
Breach Day (δ_0)	-1.307	0.036	-1.422	0.018	-1.425	0.018
Breach Day Minus 1 (δ_{-1})	0.003	0.996	-0.814	0.206	-0.816	0.207
Breach Day Minus 2 (δ_{-2})	0.640	0.211	0.862	0.124	0.859	0.123
Breach Day Minus 3 (δ_{-3})	-0.335	0.591	-0.155	0.785	-0.158	0.781
Breach Day Plus 1 (δ_1)	-0.088	0.860	-0.038	0.934	-0.042	0.927
Breach Day Plus 2 (δ_2)	0.375	0.428	0.200	0.666	0.196	0.673
Breach Day Plus 3 (δ_3)	0.136	0.772	0.006	0.990	0.007	0.988
BTC Minus 1 (ϕ)	-0.046	0.418	-0.051	0.200	-0.052	0.202
Const (c)	0.499	0.000	0.375	0.000	0.377	0.000
ARCH (α)	-	-	0.160	0.000	0.162	0.000
GARCH (β)	-	-	0.786	0.000	0.787	0.000
GJR (γ)	-	-	-	-	-0.004	0.909
Const (ω)	-	-	1.386	0.005	1.383	0.005
BIC	14,923.293	-	14,140.165	-	14,147.942	-
R^2	0.005	-	0.003	-	0.003	-
Wald Test p-value	0.141	-	0.060	-	0.061	-

Notes: Estimates of (6) are based on 2470 observations covering the January 03, 2012 to December 28, 2021 period.

Considering the parameters denoting the effects of three days before and three days after the experienced cybersecurity breach, they are all statistically insignificant at any conventional level of significance. Thus it would appear that the Bitcoin market quickly and efficiently absorbs the information relating to the cyberattacks on cryptocurrency exchanges. Parameter ϕ, which accounts for autocorrelation in Bitcoin returns, is small, negative and statistically insignificant across all three models. Finally, considering the conditional volatility parameters and observe that α and β coefficients in GARCH and GJR models are positive, statistically significant and sum up to less than one, as expected according to the existing literature, e.g. Dyhrberg (2016). Interestingly, the GJR asymmetric volatility coefficient γ is statistically insignificant at any conventional level, implying that positive and negative news shocks produce the same impact on Bitcoin volatility. In terms of model fit, the BIC criterion (Schwarz, 1978) is smallest for the GARCH model, and thus, this model is preferred over OLS and GRJ. R^2 estimates suggest that little variation of the Bitcoin return is captured by the regression equations. Finally, I present the p-values for the Wald test of overall significance of equation (6) in the last row of the table. While the OLS model's p-value exceeds 10 percent, the smaller p-values computed by the GARCH and GRJ equations suggest that equation (6) provides a better fit to the data

than a model consisting only of the intercept, at the 10 percent significance level.

Next, I consider Table 8, which presents the estimates of the model specified in equation (7) that augments the previously discussed baseline model with three control variables, namely the returns on US equities, US Dollar and gold. As before, of main interest is δ_0 which captures contemporaneous impact from recorded cybersecurity beaches of digital exchanges on Bitcoin returns. As evident from the table, δ_0 is estimated to be negative and statistically significant at the 5% level in all three cases. The estimated parameters are similar to those reported in Table 7 and now range from -1.470 (GRJ) to -1.288 (OLS). Given that the BIC criterion again favours the GARCH model, the effect of cyberattacks against cryptocurrency exchanges on Bitcoin return is estimated to be -1.459% on the day of the related cybersecurity breach. Furthermore, the market incorporates the negative news quickly with the estimated effects on the three subsequent days being statistically insignificant.

Table 8: Parameter Estimates (Added Control Variables)

	OLS		GARCH		GJR	
	Coef.	p-value	Coef.	p-value	Coef.	p-value
Breach Day (δ_0)	-1.288	0.037	-1.458	0.025	-1.470	0.023
Breach Day Minus 1 (δ_{-1})	-0.026	0.969	-0.903	0.164	-0.901	0.165
Breach Day Minus 2 (δ_{-2})	0.614	0.241	0.885	0.132	0.868	0.131
Breach Day Minus 3 (δ_{-3})	-0.409	0.524	-0.582	0.327	-0.597	0.320
Breach Day Plus 1 (δ_1)	-0.033	0.947	-0.037	0.940	-0.053	0.915
Breach Day Plus 2 (δ_2)	0.458	0.325	0.321	0.527	0.309	0.548
Breach Day Plus 3 (δ_3)	0.055	0.906	-0.156	0.744	-0.152	0.749
BTC Minus 1 (ϕ)	-0.048	0.405	-0.030	0.337	-0.032	0.325
US Equities (θ_1)	0.546	0.000	0.731	0.001	0.732	0.001
US Dollar (θ_2)	0.336	0.423	0.420	0.121	0.416	0.123
Gold (θ_3)	0.206	0.189	0.264	0.015	0.265	0.016
Const (c)	0.461	0.000	0.147	1.000	0.153	1.000
ARCH (α)	-	-	0.210	0.000	0.218	0.000
GARCH (β)	-	-	0.746	0.000	0.748	0.000
GJR (γ)	-	-	-	-	-0.020	0.669
Const (ω)	-	-	1.346	0.001	1.333	0.002
BIC	14,911.344	-	14,055.180	-	14,062.501	-
R^2	0.019	-	0.015	-	0.015	-
Wald Test p-value	0.019	-	0.007	-	0.007	-

Notes: Estimates of (7) are based on 2470 daily return observations covering the January 03, 2012 to December 28, 2021 period.

Considering the coefficients related to the three factors, I observe that the US equities market has a positive and statistically significant impact at the 1% level in all three models. The estimated coefficient varies between 0.546 in the case of OLS to 0.732 estimated by the GJR model. The effect of gold is also statistically significant in the GARCH and GRJ equations, albeit at the 5% level, while the return on the US Dollar index does not seem to influence the

return on Bitcoin at any conventional level of significance. In line with the estimates of (6), the leads and lags of the cybersecurity breach variable (I_t), as well as the lagged BTC return itself, continue to be statistically insignificant with large p-values. Finally, ARCH and GARCH coefficients are statistically significant, positive, and sum to less than one. The asymmetric γ coefficient in the GJR model is again statistically insignificant.

I complete the analysis by commenting briefly on model fit. First, the BIC criterion (Schwarz, 1978) is smallest for GARCH, implying that this specification is preferred over OLS and GRJ models. Second, despite the low estimates of R^2 which are comparable to those reported in Van Wijk (2013) and Pyo and Lee (2020) for daily return series, the Wald test of overall significance suggests that equation (7) provides a better fit to the data than a model consisting only of the intercept, at the 5 percent level.

6 Conclusions

Digital coins, such as Bitcoin, are cryptographic currencies that transact on peer-to-peer networks and allow for direct transfer of funds through online systems without an intermediary, such as a bank. Despite privacy benefits that such decentralisation provides more than 90 percent of cryptocurrency transactions still occur on centralised exchanges. This is due to convenience that organised markets offer such as easy excess, low transaction costs, and liquidity. Another key facility that centralised exchanges provide is cryptocurrency accounts where traders can store digital assets, and thus avoid the process of setting up own cryptocurrency wallets which can be technically challenging. However, the more investors enter cryptocurrency markets via centralised exchanges, the greater is the amount of client funds that is stored on online platforms run by cryptocurrency exchanges. This has led to centralised exchanges becoming attractive high-value targets to criminals, experiencing cybersecurity attacks, and having investor funds stolen.

In this thesis I investigate two types of risk associated with cyberattacks on cryptocurrency exchanges. First, I study the risk of cryptocurrency exchange closures which may result from a number of factors, one of which is weak cybersecurity programs. Historically investor funds have been either fully or partially lost when an exchange is forced to close down, and this represents a major concern that cryptocurrency traders need to take into account when choosing a digital exchange. Second, I analyse the issue of potential capital loss resulting from a price drop due to cyberattacks on cryptocurrency exchanges. This may be regarded as an indirect form of risk resulting from a market reaction to the news that a cryptocurrency exchange has been breached.

To investigate the risk of cryptocurrency exchange closures I compile a database containing eight publicly available exchange attributes on 279 cryptocurrency exchanges, 134 of which have closed since 2010. Using the collected data, I build machine learning models to predict which digital markets will remain open and which will shut down. For the prediction task I employ ten popular machine learning classifiers including i) Logistic Regression, ii) Decision Tree, iii) Random Forest, iv) Support Vector Machine, v) Multilayer Perceptron, vi) KNeighbors, vii) Naive Bayes, viii) AdaBoost, ix) ExtraTrees and x) Equally-weighted ensemble of the previous

9 classifiers. Finally, I rank the alternative algorithms according to four different measures of classification performance and identify key predictor variables.

The top three predictive models according to classification accuracy are random forest, multilayer perceptron and an equally weighted ensemble classifier. When evaluating accuracy on the training dataset the best model is the ensemble classifier reaching the accuracy of 95.9 percent. Random forest and multilayer perceptron are only marginally lower with the accuracy measures of 95.4 and 94.9, respectively. Considering out-of-sample (test dataset) accuracy, the best three models still rank in the top three position, but the random forest now performs best with 85.7 percent accuracy while multilayer perceptron, the ensemble classifier and decision tree share the second position with 83.3 percent accuracy. Three alternative measures of performance precision, recall and F1 score largely agree with the rankings provided by classification accuracy.

From the list of eight exchange characteristics, average traded volume, exchange lifetime, security audit and bug bounty program are found to be key predictors across multiple classifiers. Experiencing a previous cybersecurity breach reduces the probability of survival according to logistic regression but does not seem to impact the classification accuracy according to the decision tree and random forest classifiers. Two factor authentication and the extent of anti-money laundering regulation in the country of origin do not seem to have an impact cryptocurrency exchange closure.

Having developed a method to predict which cryptocurrency exchanges will remain active and which ones will face closure that can be utilised by traders when choosing their digital market, I next turn to analyse the impact of cybersecurity breaches of cryptocurrency exchanges on bitcoin returns. For this purpose, I use daily data covering the period January 03, 2012 – December 28, 2021, and test the hypothesis that the dates associated with cybersecurity breaches of digital exchanges experience statistically significant decreases in Bitcoin returns. The test equations account for the leads and lags of the breach date dummy variable itself, as well as other control factors such as the returns on the US equities market, US Dollar and gold. In addition, I also estimate models with and without time-varying volatility components.

The key finding here is that cybersecurity breaches of cryptocurrency markets result in a negative contemporaneous change in Bitcoin return, which is statistically significant at the 5 percent level. The impact on Bitcoin returns is estimated to range between -1.288% and -1.470%, depending on which of three alternative conditional variance models is applied, and on what control variables are included in the test equations. Interestingly the estimated coefficients on the leads and lags of the breach date indicator variable are statistically insignificant, at any conventional level of significance, indicating that the market processes the information of cyberattacks on cryptocurrency exchanges relatively quickly.

In summary, this thesis demonstrates that investors may be able to reduce their risk of exchange closures by trading on the markets that record high transaction volumes, have a long trading track record, and implement multiple security features. Results also show that a certain level of risk remains even after accounting for all the exchange characteristics that I consider in this thesis. Traders should therefore aim to stay informed of any further pertinent information, as well as consider transferring their digital assets from organised exchanges to their own cryptocurrency wallets. In addition to exercising caution when choosing which digital

exchange to transact on, cryptocurrency traders also need to be aware of potential capital losses resulting from cyberattacks on digital exchanges that may impact the price of Bitcoin itself. This type of risk may be difficult to manage as it can spread from one cryptocurrency exchange to the entire Bitcoin market.

6.1 Limitations of the Present Study

An accurate and complete dataset is a critical factor in examining and forecasting the impact of cyberattacks on cryptocurrency markets. In this thesis, I use a dataset that I have collected manually from publicly accessible sources, including various types of media, hacker forums, cryptocurrency exchange websites, and social networks. As such, there are limitations regarding the dataset, which I list below.

First, I only include cyberattacks that have been publicly announced by exchanges or through various media sources. However, there are likely other cybersecurity breaches of cryptocurrency exchanges that have not been made publicly known for multiple reasons.

Second, the breach date variable, which records the date of cyberattacks on cryptocurrency exchanges, also relies on public announcements. While in most cases, the affected exchange would detail the circumstances and timeframe of the attack, sometimes the detection of a cyberattack is delayed making it difficult to verify the actual date of the attack.

Third, when collecting data, I came across a number of cryptocurrency exchanges that did not have publicly available information on some key cybersecurity attributes, such as bug bounty programs or external security audits. In such cases, I adopted a conservative approach. I assume that exchanges did not implement security features if they did not publicly provide information. This is a reasonable assumption given that cryptocurrency investors care about the security of their funds, and digital exchanges often compete for investors based on the facilities and security features they provide. Cases of missing data were most often encountered with cryptocurrency exchanges that have closed down.

Lastly, as explained in Section 3.2, Bitcoin is traded on weekends, while markets for other assets are typically only open on weekdays. An alternative way to account for this observation is to use returns over the whole weekend for Bitcoin on the first trading day of each week. This possible extension is left for future work.

References

Azar, A. T. and El-Said, S. A. (2013). Probabilistic neural network for breast cancer classification. *Neural Computing and Applications*, 23(6):1737–1751.

Baek, C. and Elbeck, M. (2015). Bitcoins as an investment or speculative vehicle? a first look. *Applied Economics Letters*, 22(1):30–34.

Barbon, A. and Ranaldo, A. (2021). On the quality of cryptocurrency markets: Centralized versus decentralized exchanges. *arXiv preprint arXiv:2112.07386*.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327.

Borri, N., Massacci, D., Rubin, M., and Ruzzi, D. (2022). Crypto risk premia. *Available at SSRN*.

Bouri, E., Das, M., Gupta, R., and Roubaud, D. (2018). Spillovers between bitcoin and other assets during bear and bull markets. *Applied Economics*, 50(55):5935–5949.

Bouri, E., Gupta, R., Tiwari, A. K., and Roubaud, D. (2017). Does bitcoin hedge global uncertainty? evidence from wavelet-based quantile-in-quantile regressions. *Finance Research Letters*, 23:87–95.

Briere, M., Oosterlinck, K., and Szafarz, A. (2015). Virtual currency, tangible return: Portfolio diversification with bitcoin. *Journal of Asset Management*, 16(6):365–373.

Brown, M. S. and Douglass, B. (2020). An event study of the effects of cryptocurrency thefts on cryptocurrency prices. In *2020 Spring Simulation Conference (SpringSim)*, pages 1–12. IEEE.

Candila, V. (2021). Multivariate analysis of cryptocurrencies. *Econometrics*, 9(3):28.

Caporale, G. M., Kang, W.-Y., Spagnolo, F., and Spagnolo, N. (2021). Cyber-attacks, spillovers and contagion in the cryptocurrency markets. *Journal of International Financial Markets, Institutions and Money*, 74:101298.

Cheah, E.-T. and Fry, J. (2015). Speculative bubbles in bitcoin markets? an empirical investigation into the fundamental value of bitcoin. *Economics Letters*, 130:32–36.

Chen, E.-T. and Clements, A. (2007). S&P 500 implied volatility and monetary policy announcements. *Finance Research Letters*, 4(4):227–232.

Ciaian, P., Rajcaniova, M., and Kancs, d. (2016). The economics of bitcoin price formation. *Applied economics*, 48(19):1799–1815.

Civitarese, J. and Mendes, L. (2018). Bad news, technical development and cryptocurrencies stability. *Technical Development and Cryptocurrencies Stability (December 1, 2018)*.

Corbet, S., Cumming, D. J., Lucey, B. M., Peat, M., and Vigne, S. A. (2020a). The destabilising effects of cryptocurrency cybercriminality. *Economics Letters*, 191:108741.

Corbet, S., Larkin, C., Lucey, B. M., Meegan, A., and Yarovaya, L. (2020b). The impact of macroeconomic news on bitcoin returns. *The European Journal of Finance*, 26(14):1396–1416.

Dyhrberg, A. H. (2016). Bitcoin, gold and the dollar–a garch volatility analysis. *Finance Research Letters*, 16:85–92.

Feig, E. (2018). A framework for blockchain-based applications. *arXiv preprint arXiv:1803.00892*.

Gandal, N., Hamrick, J., Moore, T., and Oberman, T. (2018). Price manipulation in the bitcoin ecosystem. *Journal of Monetary Economics*, 95:86–96.

Glaser, F., Zimmermann, K., Haferkorn, M., Weber, M. C., and Siering, M. (2014). Bitcoin-asset or currency? revealing users' hidden intentions. *Revealing Users' Hidden Intentions (April 15, 2014). ECIS*.

Glosten, L. R., Jagannathan, R., and Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5):1779–1801.

Gradojevic, N. and Tsiakas, I. (2021). Volatility cascades in cryptocurrency trading. *Journal of Empirical Finance*, 62:252–265.

Guri, M. (2018). Beatcoin: Leaking private keys from air-gapped cryptocurrency wallets. *2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*, pages 1308–1316.

Hackernews (2019). A huge list of cryptocurrency thefts. [online]. Available: https://hackernoon.com/a-huge-list-of-cryptocurrency-thefts-16d6bf246389.

Hansen, P. R., Lunde, A., and Nason, J. M. (2011). The model confidence set. *Econometrica*, 79(2):453–497.

Hasanova, H., Baek, U.-j., Shin, M.-g., Cho, K., and Kim, M.-S. (2019). A survey on blockchain cybersecurity vulnerabilities and possible countermeasures. *International Journal of Network Management*, 29(2):e2060.

Hu, J., Luo, Q., and Zhang, J. (2020). The fluctuations of bitcoin price during the hacks. *International Journal of Applied Research in Management and Economics*, 3(1):10–20.

Kamiya, S., Kang, J.-K., Kim, J., Milidonis, A., and Stulz, R. M. (2021). Risk management, firm reputation, and the impact of successful cyberattacks on target firms. *Journal of Financial Economics*, 139(3):719–749.

Katsiampa, P. (2017). Volatility estimation for bitcoin: A comparison of garch models. *Economics Letters*, 158:3–6.

Klein, T., Thu, H. P., and Walther, T. (2018). Bitcoin is not the new gold–a comparison of volatility, correlation, and portfolio performance. *International Review of Financial Analysis*, 59:105–116.

Koutmos, D. (2018). Bitcoin returns and transaction activity. *Economics Letters*, 167:81–85.

Koutmos, D. (2020). Market risk and bitcoin returns. *Annals of Operations Research*, 294(1):453–477.

Kumar, A. S. and Anandarao, S. (2019). Volatility spillover in crypto-currency markets: Some evidences from garch and wavelet analysis. *Physica A: Statistical Mechanics and its Applications*, 524:448–458.

Luther, W. J. and Smith, S. S. (2020). Is bitcoin a decentralized payment mechanism? *Journal of Institutional Economics*, 16(4):433–444.

Lyócsa, Š., Molnár, P., Plíhal, T., and Širaňová, M. (2020). Impact of macroeconomic news, regulation and hacking exchange markets on the volatility of bitcoin. *Journal of Economic Dynamics and Control*, 119:103980.

Masilela, J. J., van Wyk, R. B., and Marwa, N. (2021). Assessing the variability of crypto collateral assets in secured lending on the blockchain. *Development Southern Africa*, pages 1–11.

Matkovskyy, R. (2019). Centralized and decentralized bitcoin markets: Euro vs usd vs gbp. *The Quarterly Review of Economics and Finance*, 71:270–279.

Milunovich, G. and Lee, S. A. (2022). Cryptocurrency exchanges: Predicting which markets will remain active. *Journal of forecasting*, pages 1–11. https://doi.org/10.1002/for.2846.

Moore, T. and Christin, N. (2013). Beware the middleman: Empirical analysis of bitcoin-exchange risk. In *International Conference on Financial Cryptography and Data Security*, pages 25–33. Springer.

Moore, T., Christin, N., and Szurdi, J. (2018). Revisiting the risks of bitcoin currency exchange closure. *ACM Transactions on Internet Technology (TOIT)*, 18(4):1–18.

Nakamoto, S. (2008). A peer-to-peer electronic cash system. Retreaved from https://bitcoin.org/bitcoin.pdf.

Nikkinen, J. and Sahlström, P. (2004). Impact of the federal open market committee's meetings and scheduled macroeconomic news on stock market uncertainty. *International Review of Financial Analysis*, 13(1):1–12.

Oosthoek, K. and Doerr, C. (2020). From hodl to heist: Analysis of cyber security threats to bitcoin exchanges. In *2020 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–9. IEEE.

Poletti, L. (2018). Cryptocurrency basics - 3 key characteristics and why they matter. [online]. Available: https://medium.com/the-capital/key-characteristics-of-cryptocurrency-and-why-do-they-matter-to-you-5f33e483a40f.

Pyo, S. and Lee, J. (2020). Do fomc and macroeconomic announcements affect bitcoin prices? *Finance Research Letters*, 37:101386.

Ramos, S., Pianese, F., Leach, T., and Oliveras, E. (2021). A great disturbance in the crypto: Understanding cryptocurrency returns under attacks. *Blockchain: Research and Applications*, page 100021.

RBA (2022). What are cryptocurrencies? [online]. Available: https://www.rba.gov.au/education/resources/explainers/cryptocurrencies.html.

Roubini, N. (2018). Exploring the cryptocurrency and blockchain ecosystem. *Testimony for the Hearing of the US Senate Committee on Banking, Housing and Community Affairs.*

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, pages 461–464.

Selfkey (2019). Comprehensive list of cryptocurrency exchange hacks. [online]. Available: https://selfkey.org/list-of-cryptocurrency-exchange-hacks/.

Selfkey (2020). Comprehensive list of cryptocurrency exchange hacks. Retrieved from https://selfkey.org/list-of-cryptocurrency-exchange-hacks.

Slowmist (2021). Exchange : 94 hack event(s). [online]. Available: https://hacked.slowmist.io/en/?c=Exchange.

Slowmist (2022). Exchange : 94 hack event(s). Retrieved from https://hacked.slowmist.io/en/?c=Exchange.

Stensås, A., Nygaard, M. F., Kyaw, K., and Treepongkaruna, S. (2019). Can bitcoin be a diversifier, hedge or safe haven tool? *Cogent Economics & Finance*, 7(1):1593072.

Van Wijk, D. (2013). What can be expected from the bitcoin. *Erasmus Universiteit Rotterdam*, 18.

Verdugo Yepes, C. (2011). Compliance with the aml/cft international standard: Lessons from a cross-country analysis. *IMF Working Papers*, pages 1–75.

Wallace, B. (2011). The rise and fall of bitcoin. [online]. Available: https://web.archive.org/web/20131031043919/http://www.wired.com/magazine/2011/11/mf_bitcoin.

Wilson, E. B. and Worcester, J. (1943). The determination of ld 50 and its sampling error in bio-assay. *Proceedings of the National Academy of Sciences of the United States of America*, 29(2):79.

Zhou, S. (2021). Exploring the driving forces of the bitcoin currency exchange rate dynamics: an egarch approach. *Empirical Economics*, 60(2):557–606.